

AD-A278 642





University
of Southern
California



Employing Knowledge Resources in a New Text Planner Architecture

Eduard H. Hovy, Julia Lavid
Vibhu Mittal and Cecile Paris
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292

Elisabeth Maier
Projekt KOMET, GMD-IPSI
Dolivostr, 15
D-6100 Darmstadt
Germany



INFORMATION
SCIENCES
INSTITUTE



310/822-1511

4676 Admiralty Way/Marina del Rey/California 90292-6695

Employing Knowledge Resources in a New Text Planner Architecture

**Eduard H. Hovy, Julia Lavid
Vibhu Mittal and Cecile Paris
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292**

**Elisabeth Maier
Projekt KOMET, GMD-IPSI
Dollvostr, 15
D-6100 Darmstadt
Germany**

**1992
ISI/RR-93- 372**

94-12999



94 4 28 10 11

The first author was supported in part by Rome Laboratory of the Air Force Systems Command under RL contract no. FQ7619-89-03326-0001. The second author spent a 14-month sabbatical at USC/ISI on leave from the Universidad Complutense in Madrid, Spain, and was supported by grants from the Spanish Government. Third author, a graduate student in Germany employed at the IPSI Institute in Darmstadt, Germany, spent a total of eight months working at USC/ISI and was partly supported by IPSI and partly by Rome laboratory of the Air Force Systems Command under contract no. FQ7619-89-03326-0001. The forth author, a graduate student at the University of Southern California, was supported by the Defense Advanced Research Projects Agency under contract no. NCC2-250 DABT63-91-C-0025. The fifth author was supported in part by the Defense Advanced Research Projects Agency under contract no. NCC2-250 DABT63-91-C-0025. Views and conclusions contained in this report are the authors' and should not be interpreted as representing the official opinion or policy of DARPA, Rome Laboratory, the U.S. Government, or any person or agency connected with them.

REPORT DOCUMENTATION PAGE			FORM APPROVED OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1992		3. REPORT TYPE AND DATES COVERED Research Report
4. TITLE AND SUBTITLE Employing Knowledge Resources in a New Text Planner Architecture			5. FUNDING NUMBERS FQ7619-89-03326-0001 and NCC2-250 DABT63-91-C- 0025RR-93-372	
6. AUTHOR(S) Eduard Hovy, Julia Lavid, Vibhu Mittal and Cecile Paris				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) USC INFORMATION SCIENCES INSTITUTE 4676 ADMIRALTY WAY MARINA DEL REY, CA 90292-6695			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) ROME Laboratories Griffiths AFB Rome, New York			10. SPONSORING/MONITORING AGENCY REPORT NUMBER RR-372	
11. SUPPLEMENTARY NOTES Published in 1992, Springer Verlag lecture Notes in AI #587 Aspects of Automated Language Generation editors: R. Dale, E. Hovy, D. Rosner and O. Stock				
12A. DISTRIBUTION/AVAILABILITY STATEMENT UNCLASSIFIED/UNLIMITED			12B. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) We describe in this paper a new text planner that has been designed to address several problems we had encountered in previous systems. Motivating factors include a clearer and more explicit separation of the declarative and procedural knowledge used in a text generation system as well as the identification of the distinct types of knowledge necessary to generate coherent discourse, including communicative goals, text types, schemas, discourse structure relations, and theme development patterns. This knowledge is encoded as separate resources and integrated under a flexible planning process that draws from appropriate resources whatever knowledge is needed to construct a discourse structure. We describe the resources and the planning process and illustrate the ideas with an example.				
14. SUBJECT TERMS text planning, discourse structure, computational language generation			15. NUMBER OF PAGES 19	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UNLIMITED	

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to stay within the lines to meet optical scanning requirements.

Block 1. Agency Use Only (Leave blank).

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C - Contract	PR - Project
G - Grant	TA - Task
PE - Program Element	WU - Work Unit Accession No.

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory

Block 10. Sponsoring/Monitoring Agency Report Number. (If known)

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of ...; To be published in... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."
DOE - See authorities.
NASA - See Handbook NHB 2200.2.
NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.
DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.
NASA - Leave blank.
NTIS - Leave blank.

Block 13. Abstract. Include a brief (Maximum 200 words) factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code (NTIS only).

Blocks 17.-19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.

Employing Knowledge Resources in a New Text Planner Architecture

Eduard Hovy,¹ Julia Lavid,² Elisabeth Maier,³ Vibhu Mittal,^{1,4} Cecile Paris¹

¹ USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292. U.S.A.

² Departamento de Filología Inglesa
Universidad Complutense de Madrid
280040 Madrid, Spain

³ Projekt KOMET, GMD-IPSI
Dolivostr. 15
D-6100 Darmstadt, Germany

⁴ Department of Computer Science
University of Southern California
Los Angeles, CA 90089. U.S.A.

Abstract

We describe in this paper a new text planner that has been designed to address several problems we had encountered in previous systems. Motivating factors include a clearer and more explicit separation of the declarative and procedural knowledge used in a text generation system as well as the identification of the distinct types of knowledge necessary to generate coherent discourse, such as communicative goals, text types, schemas, discourse structure relations, and theme development patterns. This knowledge is encoded as separate resources and integrated under a flexible planning process that draws from appropriate resources whatever knowledge is needed to construct a text. We describe the resources and the planning process and illustrate the ideas with an example.

1 Introduction

When generating a multisentential text, one has to perform several different types of tasks, such as selecting the material to be included, organizing this material into a coherent message, choosing words, and generating grammatical sentences. We describe in this paper a new text planner being built jointly at USC/ISI and GMD-IPSI that performs portions of the generation process. It is based on theoretical studies and experiments in text coherence, e.g., Rhetorical Structure Theory (Mann and Thompson 88), Conjunctive Relations (Martin 92), theories of discourse, e.g., (Grosz and Sidner 86, Polanyi 88),

Authors are listed in alphabetical order.

and text planning, e.g., (Hovy 88, Moore and Paris 89), advancing on those ideas and handling several new aspects of the problem.

This new text planner was designed to address several problems that we had encountered in our own systems and had observed in other, similar enterprises. An important motivation was a clearer separation of declarative and procedural knowledge in a generation system, as well as the identification of the distinct types of knowledge necessary to generate a text (for e.g., Paris and Maier 91). We had noticed in our current systems that, as the planners' plan libraries grew, the same information (e.g., requirements of use and other preconditions) had to be represented several times, and it became harder to add new plans and to modify existing plans because of their interrelationships. Furthermore, existing planning systems often mix information regarding the planning process and information necessary for linguistic realization in one single plan operator. Finally, some of the linguistic knowledge necessary to plan a text is sometimes encoded in the planner itself, rendering the process more opaque. To address these problems, we attempt to make as clear as possible the distinction between procedural and declarative information, and to identify precisely and separate out the different types of knowledge required for creating a discourse structure.

We first present the knowledge sources that we have so far identified. The planning process is then described, followed by an example of text planning.

2 Knowledge Resources Required for a Text Planner

Our text planner embodies an attempt to isolate and use some of the major knowledge resources required to plan multisentential text.² In this section, we present the major knowledge resources that we have so far identified, namely: *text types, communicative goals, schemas, discourse structure relations*, and, finally, a resource to handle *theme development and focus shift*.

In some cases, the knowledge resources embody information about the order of some planning operations. Such resources we have implemented as systemic networks; they are the discourse relations and theme patterns. In other cases, the knowledge resources provide information which the planner uses to make decisions. Such resources we have implemented as property-inheritance networks; they are the text types, communicative goals, and schemas. Both types of representation are declarative, enable us to capture inherent commonalities within the resource, and promote notational clarity and simplicity of processing.

² We consider discourse to be a structured collection of clauses. This structure is expressed by the nesting of segments of the discourse; a discourse can thus be represented as a tree structure, in which each node of the tree governs the segment (subtree) beneath it. At the top level, the discourse is governed by a single root node; at the leaves, the basic segments are single grammatical clauses. Each discourse segment has an associated purpose, which, following Grosz and Sidner (1986), we call the Discourse Segment Purpose, i.e., a communicative goal. Each adjacent pair of such purposes is related in the discourse structure by one or more discourse structure relations, as described in Section 2.4.

Each node in either type of network may contain one or more *realization operators* which indicate the effects of choosing the node, such as making additions to the discourse structure, choosing subsequent nodes to visit, setting requirements upon subsequent grammatical realization, etc. Knowledge resources co-constrain each other via these realization operators. In Section 3, we describe how the property-inheritance networks are used and the systemic networks are traversed during the planning process, and how a text structure is built during the traversal.

This planner is far from complete. Motivations for various choices have not been fully identified and several important text planning functions, such as noun phrase planning, lexical choice, lexical cohesion, and sentence structure planning are lacking altogether. We intend to investigate these issues in future work. In the remainder of this section we describe the various knowledge resources of the planner.

2.1 Text Type Hierarchy

It has long been observed that certain types of linguistic phenomena (e.g., the rhetorical structure, lexical types, grammatical features) closely reflect the genre of the text (e.g., scientific papers, financial reports). A text generation system that contains a rich set of expressive possibilities requires some representation of genres or text types in order to constrain its options, since no other resource will provide the necessary information, and the system will be unable to choose between alternative formulations.

Several text typologies have been proposed by linguists. To mention only a few: Biber (1989) identified eight basic types of texts based on statistically derived grammatical and lexical commonalities. The Washington School proposed a detailed classification of different genres of written scientific and technical English (Trimble 85), additionally pointing out typical relationships within and between rhetorical/textual units. de Beaugrande (1980) proposed a general classification of text types, also arguing that text types determine the types of discourse structure relations used.

Given its generality, we decided to base our text planner's hierarchy of text types on that of De Beaugrande, with extensions as needed to handle the text types we are addressing. The hierarchy (partially shown in Fig. 1) is represented as a property-inheritance network in the knowledge representation system Loom (MacGregor 88). Each text type in this hierarchy has associated with it the constraints it imposes on other resources, such as which communicative goals it entails, which discourse relations it favors, any appropriate grammatical constraints, etc. As a result, once a type has been established for the text to be generated, the selection of other parameters used during the generation process can be constrained appropriately (for instance, interpersonal discourse relations almost never appear in objective scientific reports, while love letters tend to contain mainly those relations). Thus our planner's predefined text types help pre-select or de-activate certain options in the generation process.

sion For	
GRA&I	<input checked="" type="checkbox"/>
TAB	<input type="checkbox"/>
owned	<input type="checkbox"/>
riation	
Distribution	
Availability Codes	
Dist	Avail and/or Special
A-1	

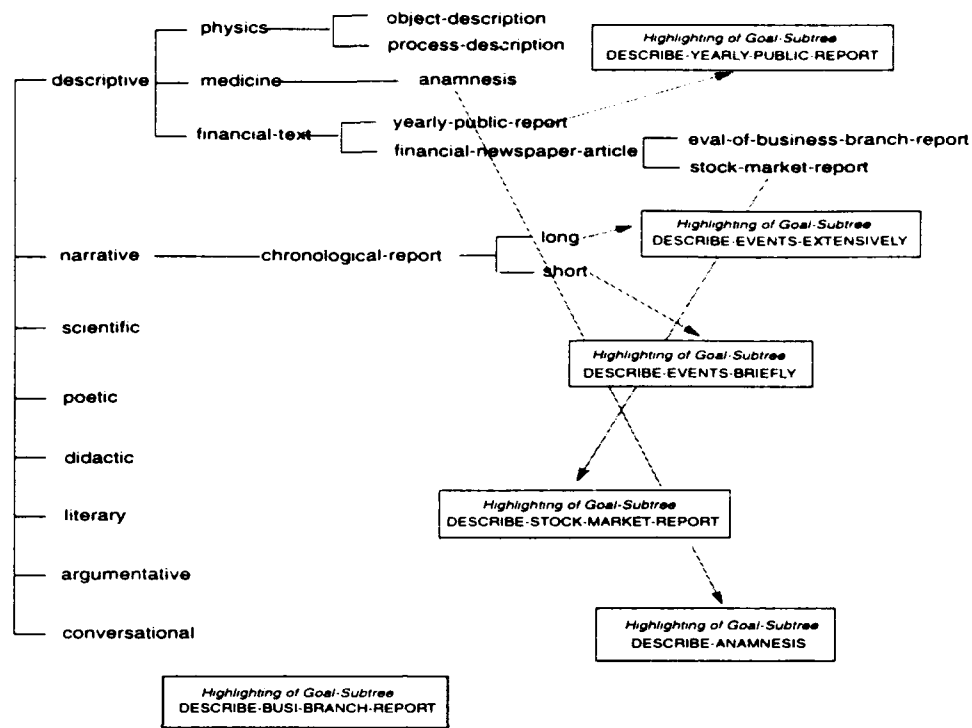


Fig. 1 A part of the Text Type Hierarchy.

2.2 Communicative Goal Hierarchy

Communicative goals have been used in many generation systems to describe the discourse purpose(s) of the speaker. Our planner contains a rudimentary taxonomization of communicative goals, starting at the topmost level with some very general goals, such as INFORM, DESCRIBE, REQUEST, and ORDER, which are eventually refined into specific goals to describe specific types of information for specific contexts (see Fig. 2). Our taxonomy, which is implemented as a property-inheritance network, resembles the one being derived from speech acts by Allen and his colleagues (e.g., Allen 91).

Planning for a communicative goal results in a discourse segment (a subtree of the discourse structure for the generated text). Each communicative goal contains one or more realization operators --- instructions for the planner to perform specific actions (see Section 3). We call our planner's lowest clause-level goals *planner primitive speech acts*; these primitive speech acts form the leaves of the discourse structure tree. During the planning process, they signal that the next step is grammatical realization.

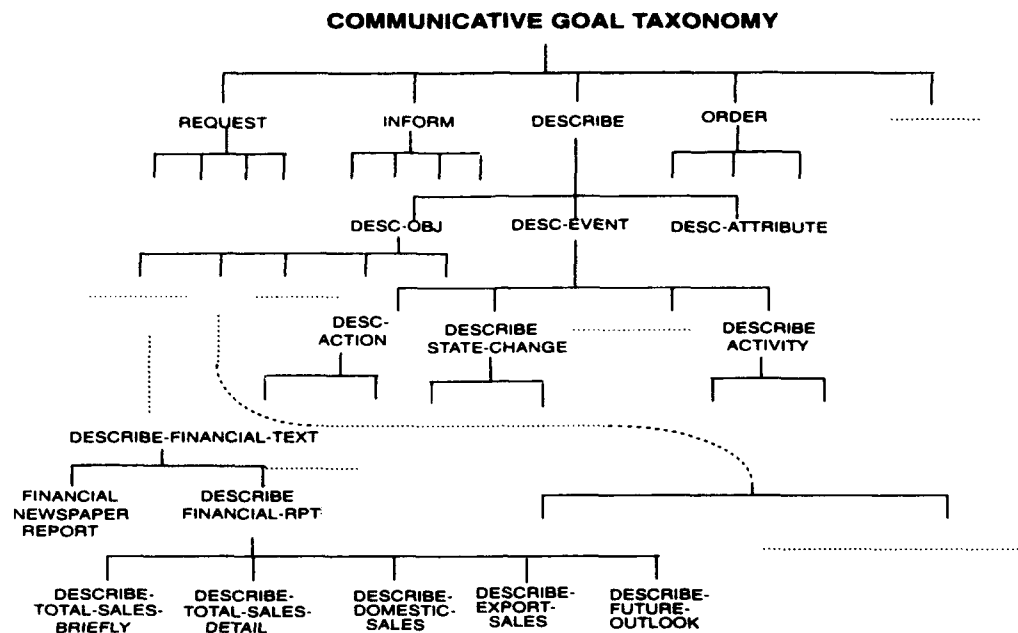


Fig. 2. A part of the Communicative Goal Hierarchy.

2.3 Schemas

In many circumstances, texts exhibit a stereotypical structure. In text planning systems, such structure is usually represented in schemas, which specify the topics of discussion that appear in the text as well as their ordering (e.g., McKeown 85). The stages of structural stereotypes can be defined at the clausal level (indicating the type of process of each sentence to be included and its position), but can equally well be defined at a more general level (indicating the sequence of general topics to be included). Linguists have proposed several schema-like approaches to model such structure: e.g., macrostructures (Van Dijk and Kintsch 83), holistic structures (Mann and Thompson 88), and the Generic Structure Potential (Halliday and Hasan 85). Recognizing the utility of such structures, we include them in our planner. We represent them as property-inheritance networks.³

A schema can be composed of several communicative goals: As an example, a

³ We believe that, in spite of the frozen nature of schemas that capture frequently occurring patterns of text structure in specific situations, there still exist rhetorical relationships among the different parts of each schema. Given sufficient knowledge, a system should be able to plan out the same text without using a schema. However,

schema to generate financial reports could contain the following communicative goals in the dictated order (see Fig. 2): (1) describe-total-sales-briefly (heading); (2) describe-total-sales-detail; (3) describe-domestic-sales; (4) describe-export-sales and (5) describe-future-outlook. We show in Section 4 how this schema is used in our planner to generate a particular text.

Just as the previous two resources co-constrain the other resources (e.g., the choice of text type can influence the selection of a schema), the instantiation of a schema can highlight or suppress different discourse relations, or the various stages of a schema can favor particular theme development patterns.

2.4 Discourse Structure Relations

Many linguists and computational linguists have studied the relationships that hold between sentences or segments of text, identifying relations that they claim need to hold in order for a text to be coherent; e.g., (Grimes 75, Mann 84, Hobbs 78, Mann and Thompson 88, Sanders *et al.* 91, Redeker 90). These relations can be used in a generation system in order to guide the selection and organization of the information to be included when other structuring guidance is lacking, such as when a schema stage calls for more material than can fit into a single clause. The necessity and use of discourse structure relations in text planners to ensure coherence has been amply discussed, as in (Hovy 88, Moore and Paris 89, Cawsey 90, Maybury 90).

Our text planner contains three networks of discourse relations, implemented as systemic networks. As a basis for the networks, we drew on several main sources: the relations defined in Rhetorical Structure Theory (Mann and Thompson 88), which were extended in Hovy's taxonomization of a collection of the relations proposed by over thirty researchers from various fields, later reorganized with Maier; see (Hovy 90, Maier and Hovy 91), and Martin's taxonomization of the conjunctive relations (Martin 92). We divided the relations into three major portions, according to the three major functions of language according to systemic linguistics (Halliday 85): semantic/ideational, interpersonal, and presentational/textual. Portions of the networks appear in Fig. 3. When organizing material, the planner is free in the general case to establish several discourse relations (typically, one for each of the major functions) between the existing discourse structure and the new piece of material; as shown in the networks, the selection of ideational, interpersonal, and textual relations is not exclusive.

As with the other resources, the discourse relation networks co-constrain the other knowledge resources, by, for example, preselecting theme patterns, posting communicative goals, or specifying aspects of grammatical realization.

lacking a complete specification of all the resources required in generation, a planner can use schemas as a useful source of 'compiled knowledge' and so avoid the need to re-derive structures over and over again.

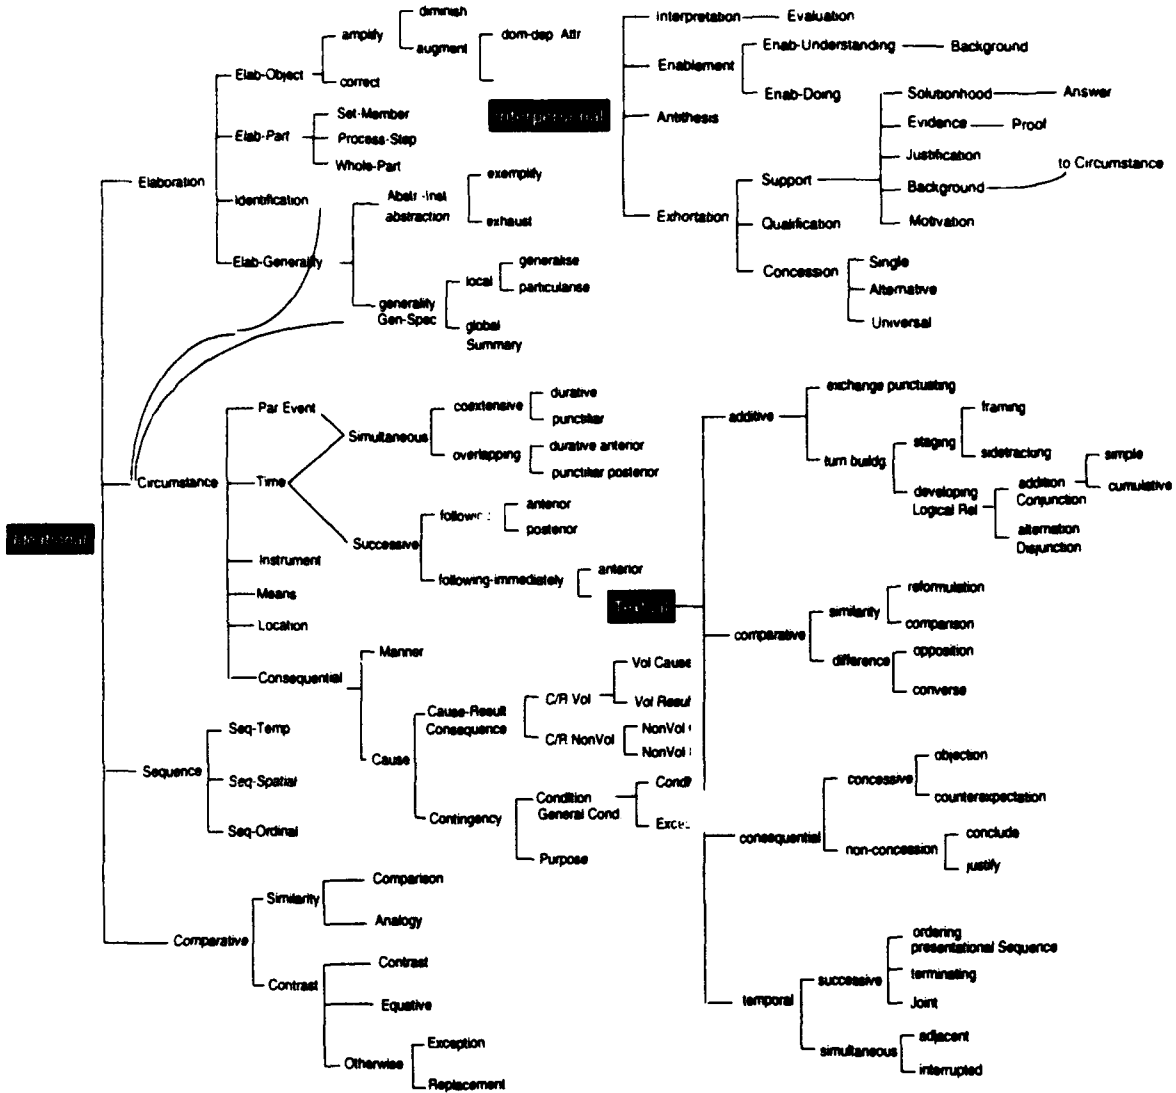


Fig. 3. Discourse structure relation networks.

2.5 Theme Development Information

Careful linguistic and computational studies have shown the need for a resource describing the potential theme developments and shifts of focus, in order to signal the introduction of a new topic of discussion and to provide its thematic relationship to previous topics (e.g., Halliday 85, Quirk *et al.* 72). Except for efforts such as Sidner (1983) and McCoy and Cheng (1988), these concerns have not been the subject of much computational work in text generation; they have generally taken the form of focus shift rules (McKeown 85, McCoy and Cheng 88, Hovy and McCoy 89). Unfortunately, these rules have usually been implemented procedurally and the complexity of the issues underlying them have not been fully studied. In the new text planner, we represent the potentialities of theme development declaratively in a systemic network (see Fig. 4).

Although the study of theme has traditionally been restricted to the sentence level, it also plays a role at the clause-complex and discourse levels. This should be taken into consideration by a text generation system. Given a text to be generated, the system must establish how theme development may proceed and how themes are to be marked in each clause. The following three concerns arise:

- the type of theme to select: following Halliday (1985), there can be three different and simultaneous themes in each clause: the ideational (or topical; expressing processes, participants, or circumstances), the interpersonal (expressing modal meanings such as probability, usuality, or opinion), and the textual (such as expressed by the continuatives "yes," "well," "oh," and similar conjunctions). The first type is semantically required.
- the theme progression pattern involved: the new theme can be the same as the theme of the previous clause; it may be part of the rheme of the previous clause; or it may be an element of what is called the "hypertheme," or general discourse segment topic (Danes 74); Note that this is similar to the focus shift rules of Sidner and McKeown.
- the linguistic degree of markedness of the theme: realization depends on the type of clause.

The motivations behind each choice follow pragmatic principles of information processing, including:

- *the topic-comment constraint* (Werth 84, Giora 88), also known as the *graded informativeness requirement*: a message is maximally effective if information which is presumed or given in the context is presented before information which is new;
- *the processibility principle* (Leech 83): a text should be constructed so that it is easy to process in real time, by placing the focus tone group at the end of the clause (the *maxim of end-focus*) and the "heavy" constituents in final position (the *maxim of end-weight*);
- *discourse relation requirements* (Mann and Thompson 88): some discourse relations have a canonical (unmarked) order of surface realization.

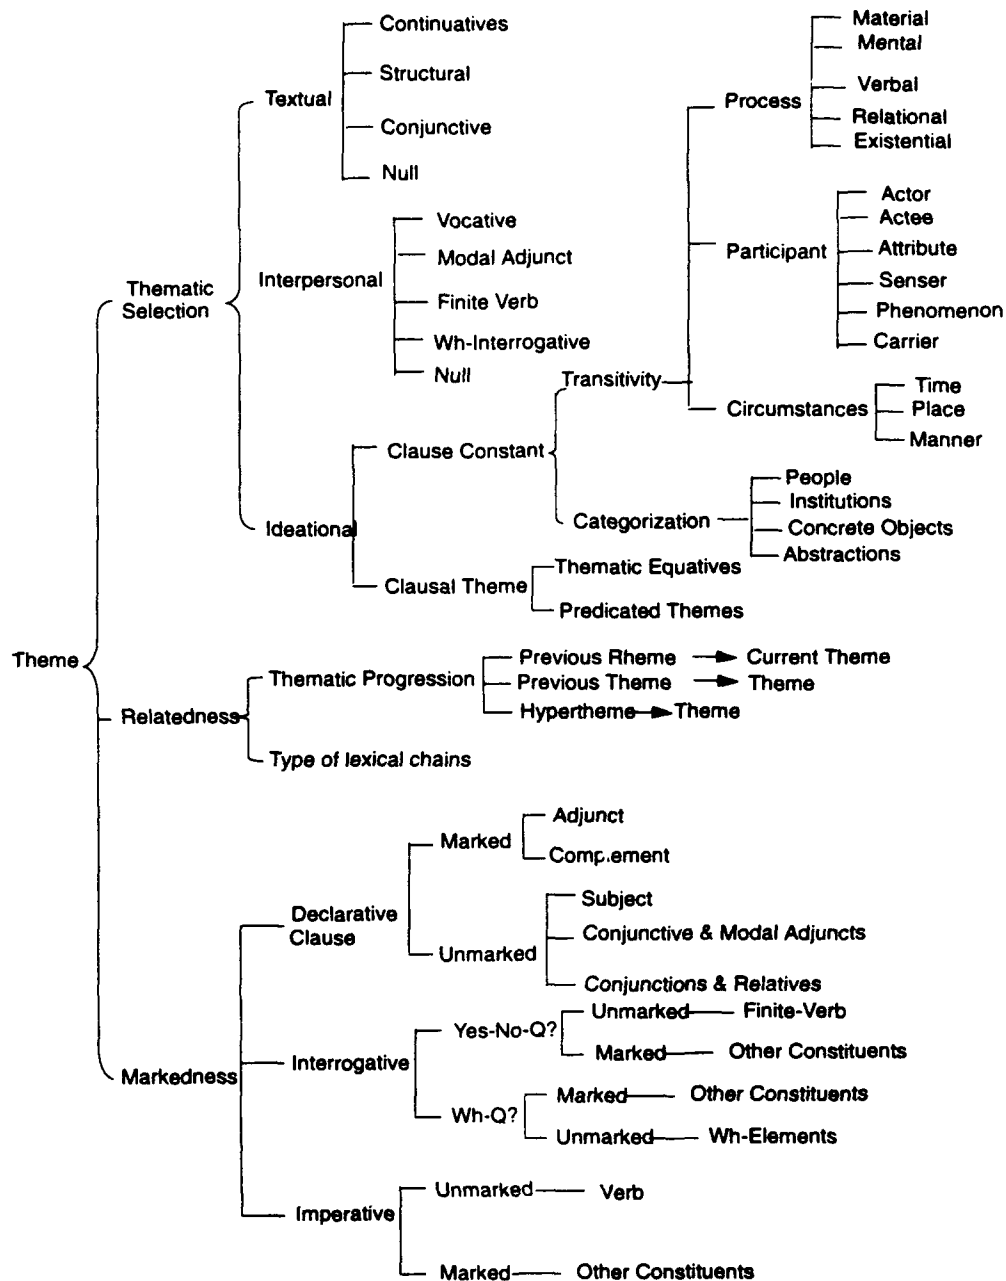


Fig. 4. A portion of the theme network.

3 The Planning Process

Planning with the networks proceeds analogously to the generation of single sentences with Penman (Penman 88, Mann 83, Mann and Matthiessen 85): in both cases, the traversal mechanism proceeds through the network, causing traversal choices to be made at nodes (systems), and building a tree-like structure as a result. We implemented the networks in Penman's internal notation so as to be able to reuse some of its traversal code.

Associated with each node in the networks is an inquiry function which queries the environment in order to determine which branch to follow, and a set of realization operators that instruct the planner what to do next.

The planning operation is very simple. After an initial setup phase, the system simply executes a basic planning cycle over and over again until planning is complete. In the setup phase, the user activates the planner with a communicative goal, as described in Section 2.2, which causes the selection of a desired text type, and is then posted on the goal stack and simultaneously on the Discourse Structure Tree. Then the basic planning cycle begins. Essentially, this cycle proceeds as follows: First, the planner checks whether there is a realization on the agenda. If so, it performs the realization by applying its action to its parameters. If there are no realizations left, the planner checks whether there is a discourse goal on the goal-stack. If there is, the planner finds the realizations associated with the goal and loads them onto the agenda; if no discourse goals remain, the planning is done.

Clearly, the action of the system lies in the realizations. Each realization is an instruction to be performed. At present, the system uses the following realizations:

1. (ACTIVATE-SCHEMA <schema-name>): Find the schema and load its realizations onto the agenda.
2. (ADD-TO-D-STRUC <goal> <concept> <parentpos>): Add the given communicative goal into the discourse structure tree at the given position.
3. (CHANGE-HYPERTHEME <chain-of-roles>): Change the topic under discussion to the filler of the given chain of roles, starting from the current topic.
4. (HIGHLIGHT-COMM-GOALS <goals>): Highlight the given goals so that only they will be considered for future planning.
5. (HIGHLIGHT-RELATION <relations>): Start traversal of the discourse relations network(s) at the given relations, using the current topic of discussion.
6. (BLOCK-RELATION <relations>): Mark the given discourse structure relations so that they cannot be traversed for the remainder of the current sentence.
7. (PREFER-THEME <concept-role>): Add instructions for the realization component that the given role of the topic under discussion should be thematized in the clause.
8. (SET-MACROTHEME <concept>): Change the overall topic of discussion.

9. (SET-UP-DISCOURSE-GOAL <goal>): Activate the given goal: load it onto the goal stack and into the discourse structure tree at the current growth point and add its realizations to the agenda.
10. (TRAV-ONE-NETWORK-NODE <node-name>): Locate the given node in the knowledge resource networks, apply its inquiry function, record the result (the inquiry choice), and load the realizations associated with the result onto the agenda.

4 An Example of the Planner in Action

In this section we provide a brief trace in order to show how the various linguistic resources interact to guide the construction of the discourse structure. As example we take a fragment of a text from a bank's annual report:

- (1) Declines in Total Sales of the Swiss Cheese Union
- (2) In the business year 1986/87 (ending July 31), the 40 cheese trading firms associated in the Swiss Cheese Union sold 79,035 tons of cheese altogether, (3) equal to a 2.6% decline. (4) Domestic sales of table cheeses enjoyed a relatively positive trend, (5) with Swiss households buying 22,100 tons of their preferred cheeses, (6) a gain of 3.9% from one year earlier.

We represented the semantic information in this text in the Loom knowledge representation system.

Given the communicative goal GENERATE-YEARLY-PUBLIC-REPORT and the topic of discussion CHEESE-UNION-SALES-86, the schema mentioned in Section 2.3 is activated, and the planner goes through the stages indicated in the schema. Let us assume now that the first two clauses -- the headline and the first clause -- have already been generated. The state of the discourse structure appears in Fig. 5.

In generating (2), the goal DESCRIBE-TOTAL-SALES-DETAIL is popped off the goal stack, and its definition in the goal hierarchy is checked. This new information is now incorporated in the discourse structure, and the realization statements for the goal are loaded onto the agenda, including the one to highlight the discourse structure relation interpretation.

In the next cycle, the planner checks the agenda and finds the just-loaded realization. It performs the realization by highlighting interpretation in the interpersonal relations network, which causes the planner to check whether any topic material with that relation to the current topic of discussion can be brought into the discourse. This check is performed by an inquiry function that accesses the planner environment with a question that can be paraphrased as:

Interpretation-Q-Code:

"Was a numerical value mentioned in the last proposition and can it be expressed in relation to other values?"

From the information about the topic (as contained in the knowledge representation system), a possible candidate for such a relation is the value of the role weight-ascription. The inquiry code retrieves a role and a value which




Fig. 5. Snapshot of the text planner state.

fulfills the above condition: the role *weight-change-relative* represents the weight ascription relative to that of the preceding year. The relevant segment of the domain model appears in Fig. 6.

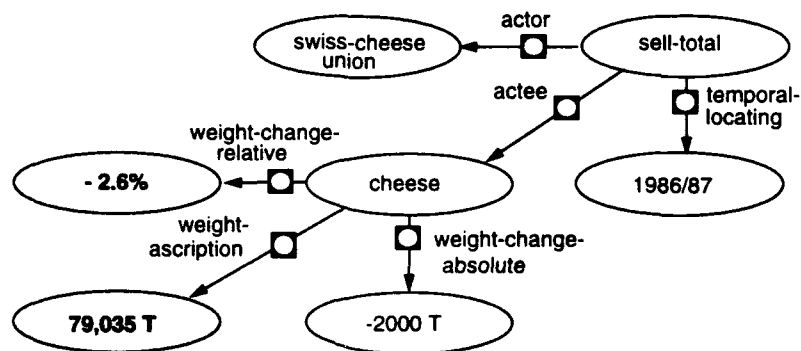


Fig. 6. Fragment of the Domain Model.

The successful finding of this material signals the applicability of the relation interpretation. The planner then activates the realization statements associated with this relation, in this case:

1. **knowledge selection:** Each relation contains specifications of the material it relates. The realization associated with interpretation selects both the absolute and the relative ascriptions for the weight, and calls for incorporating this new information.
2. **discourse structure growth:** This realization calls for the addition of the interpretation relation at the current growth point in the discourse structure tree of Fig. 5.
3. **theme determination:** This realization calls for traversal of the theme network in order to determine the thematization pattern of the new clause or clauses.
4. **operations on relations:** To prevent the repetitive use of the interpretation relation, which would lead to a monotonous text, this realization calls for interpretation to be blocked for further use until the end of the next sentence.

The planner loads these four realizations onto the agenda and thereby completes its cycle.

In the next cycle, the planner runs the knowledge selection realization listed above and selects the new material. In the following cycle, it adds the interpretation relation to the discourse structure tree. This relation indicates that clause (3) is related to clause (2) by the rhetorical relation interpretation. The resulting form of the discourse structure after these realizations appears in Fig. 7. In the following cycles, the remaining realization statements are executed.




Fig. 7. Discourse structure after the new relation has been planned.

We cannot describe the rest of the planning in detail here. In essence, the planning cycle keeps repeating, first handling all the realizations on the agenda and then all the goals on the goal stack, until no more remain.

5 Conclusion

We have presented in this paper a new text planner currently being developed jointly at USC/ISI and GMD-IPSI. It is based on the idea that the linguistic resources needed to generate coherent text (as well as their interrelationships) should be represented explicitly, separately, and distinct from the procedural knowledge required for text planning.

We do not claim to have identified all the knowledge sources required to produce coherent discourse. We are only starting to study the issues of lexical choice and lexical cohesion in depth (Wanner and Maier 91), and the planning of noun groups (and referring expressions in general), the problem of sentence delimitation are all unaddressed in our planner. We do, however, believe that the architecture of our planner lends itself well to the incorporation of additional knowledge resources when they become available. The representational power of systemic networks --- interlocking options that capture the potentialities of expression --- and the clear and simple planning cycle offer, we hope, sufficient scaffolding for the needs of text planning of the future.

Other issues of course remain to be addressed as well. For example, will we be able to keep the complexity of the procedural aspect of the system to a minimum? As the inquiry operators are implemented procedurally, we clearly need to limit how much can be done in these operators, or, at least, we need to apply clear constraints to be used when writing these operators lest their complexity becomes such that they are hard to understand and modify. We also need to study how the architecture addresses some of the concerns raised by some researchers regarding top-down planners (e.g., Suthers (1991) and Mooney, Carberry and McCoy (1990)). It is too early to discuss these issues. We will address them once the system has been more fully implemented and used to generate a greater variety of texts.

6 Acknowledgments

This planner was designed mainly during a period when both Lavid and Maier were visiting USC/ISI. In addition, many people have contributed, directly or indirectly, to the ideas embodied in this planner. While it is impossible to thank them all, we wish to acknowledge Beatrice Cahour (INRIA, France), Giuseppe Carenini (IRST, Italy) and Richard Whitney (USC/ISI) for their participation in the design.

In addition, we wish to thank the following sponsors: Hovy was supported by the Rome Air Development Center (now Rome Laboratories) under RADC contract FQ7619-89-03326-0001. Paris and Mittal were supported by NASA-Ames grant NCC 2-520 and under DARPA contract DABT63-91-C-0025. During

her one-year visit to USC/ISI, Lavid was supported by the Spanish Ministry of Education under the PROEX program. During the 5 months Maier visited USC/ISI, she was supported under RADC contract FQ7619-89-03326-0001.

References

- Allen, J.F. Discourse Structure in the TRAINS Project. In *Proceedings of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann, San Mateo, 1991.
- Biber D. A Typology of English Texts. *Linguistics*, 27, 1989, pp. 3-43.
- Cawsey, A. Generating Communicative Discourse. In *Current Research in Natural Language Generation*, Dale, R., Mellish, C., and Zock, M. (eds), Academic Press, Boston, 1990, pp. 75-102.
- Clark, H.H. and Clark, E.V. Semantic distinctions and memory for complex sentences. *Quarterly Journal of Experimental Psychology*, 20, 1968, pp. 129-138.
- Danes, F. Functional sentence perspective and the organization of the text. In F. Danes (ed.) *Papers on Functional Sentence Perspective*. Mouton, The Hague, 1974.
- de Beaugrande, R. *Text, Discourse and Process*. Ablex Publishing Corporation, Norwood, N.J., 1980.
- Giora, R. On the Informativeness Requirement. *Journal of Pragmatics* 12, 1988, pp. 547-565.
- Grimes, J.E. *The Thread of Discourse*. Mouton, The Hague 1975.
- Grosz, B.J. and Sidner, C.L. Attention, Intentions, and the Structure of Discourse. In *Journal of Computational Linguistics* 12(3), 1986, pp. 175-204.
- Halliday, M.A.K. *An Introduction to Functional Grammar*. Edward Arnold Press, Baltimore, 1985.
- Halliday, M.A.K. and Hasan, R. *Language, Context and Text: Aspects of Language in a Social-Semiotic Perspective*. Deakin University Press, 1985.
- Hobbs, J.R. Why is Discourse Coherent? Technical Note no.176, SRI International, Menlo Park CA, 1978.
- Hovy, E.H. Planning coherent multisentence text. *Proceedings of the ACL*, Buffalo, 1988.
- Hovy, E.H. Parsimonious and Profligate Approaches to the Question of Discourse Structure Relations. In *Proceedings of the 5th International Workshop on Text Generation*, Pittsburgh, 1990.
- Hovy, E.H. and McCoy, K.F. Focusing your RST: A step toward generating coherent multisentential text. In *Proceedings of the 11th Cognitive Science Conference*, Ann Arbor, 1989, pp. 667-674.
- Leech, G.N. *Principles of Pragmatics*. Longman, London, 1983.
- MacGregor, R. A Deductive Pattern Matcher. In *Proceedings of the 6th National Conference on Artificial Intelligence*, 1988, pp. 696-701.
- Maier, E. and Hovy, E.H. A Metafunctionally Motivated Taxonomy for Discourse Structure Relations. In *Proceedings of the 3rd European Workshop on Language Generation*, Judenstein, Austria, March 1991, pp. 38-45.
- Mann, W.C. and Matthiessen, C.M.I.M. Nigel: A Systemic Grammar for Text Generation. Research Report RR-83-105, USC/ISI. 1983.
- Mann, W.C. An Overview of the Penman Text Generation System", Research Report ISI/RR-83-114, USC/ISI, 1983.
- Mann, W.C. Discourse Structures for Text Generation. In *Proceedings of the COLING Conference*, Stanford, 1984.

- Mann, W.C. and Matthiessen, C.M.I.M. Nigel: A Systemic Grammar for Text Generation. In R.Benson and J.Greaves, (eds), *Systemic Perspectives on Discourse: Selected Papers Papers from the 9th International Systemics Workshop*. Ablex, London, 1985.
- Mann, W.C. and Matthiessen, C.M.I.M. and Thompson, S.A. Rhetorical Structure Theory and Text Analysis. Research Report ISI/RR-89-242, USC/ISI, 1989.
- Mann, W.C. and Thompson, S.A. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text* 8(3), 1988, pp. 243-281. Also available as USC/ISI Research Report RR-87-190.
- Martin, J.R. 1992. *English Text: System and Structure*. Benjamins, Amsterdam. Forthcoming.
- Maybury, M.T. *Planning Multisentential English Text Using Communicative Acts*. Ph.D.dissertation, Cambridge University, 1990. Also available as RADC Technical Report 90-411, 1990.
- McCoy, K.F. and Cheng, J. Focus of attention: Constraining what can be said next. In *Natural Language in Artificial Intelligence and Computational Linguistics*, Paris, C.L., Swartout, W.R. and Mann, W.C.(eds), Kluwer, Boston, 1990, pp. 103-124.
- McDonald, D.D. and Pustejovsky, J. Description-directed natural language generation. In *Proceedings of the 9th IJCAI*, Los Angeles, 1985.
- McKeown, K.R. *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press, Cambridge, 1985.
- Mooney, D.J., Carberry, S. and McCoy, K.F. The Basic Block Model of Extended Explanations. In *Proceedings of the AAAI Workshop on Explanations*, Boston, 1990.
- Moore, J.D. and Paris, C.L. Planning Text for Advisory Dialogues. In *Proceedings of the 27th ACL Conference*, Vancouver, 1989, pp. 67-75.
- Paris, C.L. and Maier, E. Knowledge Sources or Decisions? *Proceedings of the Workshop on Decisions in Natural Language Generation*, IJCAI 1991, Sydney.
- The Penman Documentation*. 5 unpublished volumes, USC/ISI.
- Polanyi, L. A Formal Model of the Structure of Discourse. *Journal of Pragmatics*, 12, 1988, pp. 601-638.
- Prideaux, G.D. Text data as evidence for language processing principles: The grammar of ordered events. In *Language Sciences*, 11, 1989, pp. 27-42.
- Prideaux, G.D. Syntactic form and textual rhetoric: The cognitive basis for certain pragmatic principles. *Journal of Pragmatics*, 16, 1991, pp. 113-129.
- Quirk, R., Greenbaum, S., Leech, G. and Svartvik, J. *A Grammar of Contemporary English*. Seminar Press, New York, 1972.
- Redeker, G. Ideational and Pragmatic Markers of Discourse Structure. *Journal of Pragmatics* 14, 1990, pp. 367-381.
- Sanders, T.J.M., Spooren, W.P.M.S., and Noordman, L.G.M. Towards a Taxonomy of Coherence Relations. In *Discourse Processes*, 1992 (forthcoming)
- Sidner, C.L. Focusing and Discourse. *Discourse Processes* 6, 1983, pp. 107-130.
- Suthers, D.D. A Task Appropriate Hybrid Architecture for Explanation. *Computational Intelligence*, 7(4), 1991.
- Trimble, L. *English for Science and Technology. A Discourse Approach*. Cambridge University Press, Cambridge, 1985.
- Van Dijk, T.A. and Kintsch, W. *Strategies of Discourse Comprehension*. Academic Press, New York, 1983.
- Wanner, L. and Maier E. Lexical Choice as an Integrated Component of Situated Text Planning, *Proceedings of the 3rd European Natural Language Generation Workshop*, Innsbruck, March, 1991.
- Werth, P. *Focus, Coherence and Emphasis*. Croom Helm, London. 1984.

This article was processed using the \LaTeX macro package with LMAMULT style